P515017.PDF [Page: 3 of 33]

Copy No. 10

DEFENCE R&D DÉFENSE

A Probabilistic Theory for the Design of Optimal Linear Discriminators for the Automated Detection of Objects in Sidescan Sonar Images

Ronald T. Kessel
DREA

DEFENCE RESEARCH ESTABLISHMENT ATLANTIC

Technical Memorandum

DREA TM 2000-021

March 2000

National Defence Défense nationale

Canada

P515017.PDF [Page 1 of 33]

A Probabilistic Theory for the Design of Optimal Linear Discriminators for the Automated Detection of Objects in Sidescan Sonar Images

Ronald T. Kessel
DREA

Defence Research Establishment Atlantic

Technical Memorandum

DREA TM 2000-021

March 2000

Author

Ronald T. Kessel

Approved by

R. Kuwahara

Head, Electromagnetics Section

Approved for release by

Walter E. Eliis

Chair, DREA Document Review Panel

© Her Majesty the Queen as represented by the Minister of National Defence, 2000

© Sa majesté la reine, représentée par le ministre de la Défense nationale, 2000

102-02589/

Abstract

(U) Computerized pattern recognition can be used to help a sonar operator locate underwater objects in sidescan sonar images. The theory behind several linear discriminators is presented here with a view to improving this automated detection. The discriminators are optimal insofar as they maximize the detection performance as defined under the Neyman-Pearson design criteria, with adjustments made to those criteria to suit the prior knowledge of both the objects sought and the local seafloor clutter. The emphasis throughout is on sea mine detection in naval operations. The theory gives practical insight and direction for the mine detection problem, showing, for instance, 1) what kind of data should be extracted from target and clutter image libraries to get optimal detection performance; 2) that the matched filter, favored for its simplicity, is just one of several optimal linear discriminators resulting in this case when nothing is assumed about the local seafloor clutter; and 3) that prior de-meaning of images will in some cases improve detection performance.

(U) La reconnaissance informatisée de formes peut être utilisée pour aider l'opérateur sonar à localiser des objets sous-marins sur des images de sonars à balayage latéral. Nous présentons ici la théorie qui s'applique à plusieurs discriminateurs linéaires dans un souci d'améliorer la détection informatisée. Les discriminateurs sont optimaux en ce sens qu'ils maximisent la performance de détection définie en vertu des critères de conception de Neyman-Pearson; des ajustements ont été faits à ces critères pour tenir compte de la connaissance préalable des objets cherchés et des bruits parasites du fond marin en question. En tout temps, l'accent est mis sur la détection de mines marines dans les opérations navales. La théorie offre un aperçu et une orientation pratiques pour le problème de détection de mines, indiquant, par exemple, 1) quelle sorte de données sur les cibles et les bruits parasites devraient être extraites des bibliothèques numériques pour obtenir une performance optimale de détection; 2) que le filtre adapté, préféré en raison de sa simplicité, n'est qu'un des nombreux discriminateurs linéaires optimaux utilisé lorsqu'on ne connaît rien à propos des bruits parasites du fond marin en question; et 3) que la soustraction de la moyenne des images va, dans certains cas, améliorer la performance de détection.

DREA TM 2000-021

A probabilistic theory for the design of optimal linear discriminators for the automated detection of objects in sidescan sonar images

Ronald T. Kessel

EXECUTIVE SUMMARY

Introduction

When sidescan sonars are used in mine hunting, the sonar operator must recognize mines against a background of highly variable clutter that depends on the quality of the sonar, its altitude above the seafloor, and on the nature of the seafloor—whether it is sand, rock, silt, sloping and so forth. Computer image processing is often used to help the operator identify mines and reject false targets. Here we consider the fundamental challenge for such processing, namely: What computer processing gives the best performance?

Principal Results

The optimal (linear) detector is formulated mathematically assuming different levels of prior information about the mines sought and the local seafloor clutter.

Significance of the Results

The theory of optimal detectors gives insight and direction for mine detection algorithms, showing how to achieve the best performance for a given prior knowledge, while clarifying what additional data, extracted from a library of actual sonar images ("mug shots") of known targets and clutter, would improve detection still further. This theoretical framework may prove to be an important advance for the Mine Countermeasures community. The optimal discriminator can also be applied more widely in non-acoustic pattern recognition problems, such as buried mine detection by electromagnetic scattering, by substituting time-sampled electromagnetic signatures for image intensities.

Future Plans

Research into optimal detection continues with a view to providing automatic detection algorithms to be used in the sidescan sonar stations on board the Canadian Forces Kingston Class vessels. This research includes compiling a library of mine and non-mine classes from actual sonar data, developing adaptive optimal discriminators that are continually apprised of the local seafloor clutter, and testing the effects of various kinds of image preprocessing on detection (such as image compression, de-noising, and non-linear changes of intensity).

CRDA TM 2000-021

Théorie probabiliste pour la conception de discriminateurs linéaires optimaux pour la détection informatisée d'objets sur des images de sonars à balayage latéral

Ronald T. Kessel

RÉSUMÉ

Introduction

Lorsque les sonars à balayage latéral sont utilisés pour la recherche de mines, l'opérateur sonar doit reconnaître les mines parmi de nombreux bruits parasites très variables qui dépendent de la qualité du sonar, de son altitude par rapport au fond marin et de la nature du fond marin, qu'il soit fait de sable, de roche ou de limon, qu'il soit incliné, etc. Le traitement informatisé des images est souvent utilisé pour aider l'opérateur à identifier les mines et à rejeter les cibles fausses. Le recours à un tel procédé soulève une question importante, à savoir : quel traitement informatique offre la meilleure performance ?

Résultats principaux

Le détecteur optimal (linéaire) est mathématiquement configuré d'après les renseignements connus au sujet des mines cherchées et des bruits parasites du fond marin en question.

Signification des résultats

La théorie du détecteur optimal offre un aperçu et une orientation pour les algorithmes de détection de mines, indiquant comment atteindre la meilleure performance pour une connaissance préalable donnée, tout en clarifiant quelles données additionnelles, [illegible] d'une bibliothèque d'images réelles de sonars (" photos signalétiques ") de cibles et de bruits parasites connus, amélioreraient davantage la détection. Ce cadre théorique peut s'avérer une étape importante pour la communauté de lutte contre les mines. Le discriminateur optimal peut aussi s'appliquer de manière plus générale aux problèmes de reconnaissance des formes non acoustiques, comme la détection de mines enterrées par diffusion électromagnétique, en substituant des signatures électromagnétiques à échantillonnage temporel aux intensités d'image.

Projets futurs

La recherche au niveau de la détection optimale se poursuit dans un souci d'offrir des algorithmes de détection automatique à utiliser dans les stations de sonar à balayage latéral à bord des navires de la classe Kingston des Forces canadiennes. Cette recherche comprend la création d'une bibliothèque dans laquelle les données réelles des sonars seront divisées en deux catégories : les données liées aux mines et celles qui ne le sont pas. Elle comprend aussi le développement de discriminateurs optimaux adaptatifs qui détectent de façon continue des bruits parasites du fond marin en question, et des essais

visant à déterminer les effets des divers types de prétraitement d'images sur la détection (comme la compression d'images, l'élimination des bruits et les changements non linéaires d'intensité).

Contents

1 Introduction 2

2 Linear discriminators 5
2.1 Performance probabilities 6
2.2 Prior information 7

3 Unapprised discriminators of the first and second kind 9
3.1 Matched filter 10
3.2 Mean and energy detector 11

4 Apprised discriminators of the first and second kind 12
4.1 Numerical solution 13
4.2 Correlation and matched filter 14
4.3 Improved matched filter performance by de-meaning 15

5 Conclusions 16

Chapter 1

Introduction

Computer algorithms for the detection of objects in sidescan sonar images make use of image processing for feature enhancement, and then pattern recognition to identify those distinctive features amid seafloor clutter. In naval operations the objects sought may be sea mines endangering ship traffic during times of hostilities. That is the context assumed here, although the theory applies more generally in the search for other objects as well. Here we focus particularly on the problem of pattern recognition, addressing what may be its most fundamental question: What discriminator gives the best detection performance?

Optimal discriminators have been treated at length in the literature [1], [2], [3], albeit very generally, apart from either image processing or mine hunting, and apparently only for recognition problems in which the cost of errors and the prior class probabilities are available, making it possible to minimize the operational costs (or Bayesian risks) directly. This cannot be done for mine hunting where the costs and prior probabilities of encountering a mine are usually unknown. Alternative performance specifications such as the Neyman-Pearson criteria [4] are therefore required.

As part of the detector's design specifications, the Neyman-Pearson criteria ordinarily prescribe a maximum allowable probability of false alarm, while maximizing the probability of detection under that constraint [5]. Variants of these criteria are their reverse, to prescribe a minimum allowable probability of detection while minimizing the probability of false alarm; or more simply, to maximize the probability of detection while ignoring the probability of false alarm; or finally, what is rarely done in practice, to minimize the probability of false alarm while ignoring the probability of detection. The

Prior Knowledge	Optimization criteria	Optimal Discriminator
of targets & clutter	prescribe P_0; maximize P_1	*apprised* of the *first* kind
of targets & clutter	prescribe P_1; minimize P_0	*apprised* of the *second* kind
of targets only	ignore P_0; maximizeP_1	*unapprised* of the *first* kind
of clutter only	ignore P_1; minimize P_0	*unapprised* of the *second* kind

Table 1.1: The prior knowledge of the target and clutter classes determines the optimization criterion, which in turn determines the form of the optimal discriminator. P_0 is the probability of false alarm, and P_1 is the probability of detection. The terms "apprised" and "unapprised" refer to the state of knowledge, whether the discriminator is fully apprised of both classes or not. Unapprised discriminators are considered in Chapter (2); apprised in Chapter (3). The terms "first" and "second kind" identify which probability is optimized: P_1 being maximized in discriminators of the first kind, and P_0 minimized in the discriminators of the second.

situation is summarized in Table (1.1).

One would presumably never ignore one or another probability unless forced to by the lack of prior information about one object class or the other—targets or clutter—making it impossible to account for both at once in the design of the detector. For lack of a better name, optimal discriminators that are fully apprised of both targets and clutter shall be called *apprised* discriminators, and those unapprised of one class or the other shall be called *unapprised* discriminators (see column three in Table (1.1)).

In this paper we derive the optimal, data based, linear discriminator for each case listed in Table (1.1). "Linear" here means that the target/non-target decision is made by thresholding the output of a linear operator applied to the grey-scale pixel intensities in the scanned image. Justification for the use of linear discriminators in the sonar detection problem is given in [6]. "Data based" means that the discriminator is derived from actual target and clutter sonar images collected under controlled experiments, or synthesized using computer simulations or physical models. "Optimal" means that the discriminator is derived here using the calculus of probabilities to maximize the performance for a given state of knowledge.

Although data based, the mathematical forms of the optimal discriminators are very general and therefore give practical insight and direction for mine detection. In particular, they show how to achieve the best performance for a given prior knowledge, while also clarifying what additional information

should be compiled and extracted from target and clutter image libraries to improve matters further. It is also shown that the *matched filter* now in wide use [7], [8], [6], [9] is a simplified version of an *unapprised* discriminator resulting when nothing is assumed about the seafloor clutter (Section (3.1)), and when prior *de-meaning* all images before processing can improve detection performance (Section (4.3)).

We begin with a review of linear discriminators.

Chapter 2

Linear discriminators

The simplest discriminator for sonar is the energy detector, in which a detection is registered whenever the return echo exceeds a certain threshold [10], [11]. In mine hunting the high energy may be the echo from a mine-like object on the seafloor. But a modern side scan sonar creates a plan-view image of the seafloor which provides much more information about the object, especially about its size and shape. The next step in detector complexity, then, is to perform a linear operation on a portion of the sonar image in a way that yields a discriminant c, with a detection registered whenever c exceeds a given threshold c_T,

$$c > c_T \text{ for detection.} \tag{2.1}$$

If the portion of the image in question is a rectangle R pixels high and C pixels wide, for instance, and the echo intensity at each pixel inside that window is M_{ij}, for $i = 1 \ldots R$ and $j = 1 \ldots C$, then the most general linear operation is

$$c = \sum_{i=1}^{R} \sum_{j=1}^{C} F_{ij} M_{ij}, \tag{2.2}$$

where F_{ij} are the discriminator coefficients to be determined for optimal performance.

To this end it is convenient to write the matrices M and F as their one dimensional isomorphic counterparts, vectors $\mathbf{m}$ and $\mathbf{f}$, by placing the matrix rows end to end

$$\mathbf{m} = \left[\begin{array}{cccccccccc} M_{1,1} & M_{1,2} & M_{1,3} & \cdots & M_{1,C} & M_{2,1} & M_{2,2} & M_{2,3} & \cdots & M_{R,C} \end{array} \right]^T, \tag{2.3}$$

$$\mathbf{f} = \left[\begin{array}{cccccccccc} F_{1,1} & F_{1,2} & F_{1,3} & \cdots & F_{1,C} & F_{2,1} & F_{2,2} & F_{2,3} & \cdots & F_{R,C} \end{array} \right]^T, \tag{2.4}$$

in which case (2.2) simply becomes

$$c = \mathbf{m} \cdot \mathbf{f}. \tag{2.5}$$

2.1 Performance probabilities

As mentioned, the Neyman-Pearson criteria are used when the cost and prior probability of encountering a mine are not known. The performance is assessed instead in terms of the conditional probabilities of detection and false alarm, which play a decisive role on the operational risk but are not a complete formulation of it. The *conditional probability of detection* P_1 is the probability that (2.1) is true given an image of a mine, and the *conditional probability of false alarm* P_0 is the probability that (2.1) is true given an image without a mine,

$$P_{0,1} = P\left(c > c_T \mid \text{mine present} = 0 \text{ (false)}, 1 \text{ (true)}\right). \tag{2.6}$$

The discriminant c is a random variable because signal noise, seafloor clutter, and uncertainties in the mine's orientation, make it impossible to predict its exact value. c therefore has some conditional probability density $p_1(c)$ given a mine, and another $p_0(c)$ given no mine, whereby

$$P_i = \int_{c_T}^{\infty} p_i(c)\, dc, \tag{2.7}$$

in which $i = 0$ for the non-mine class, or 1 for the mine class. The situation is shown graphically in Fig.(1).

If an optimal linear discriminator $\mathbf{f}$ exists, then P_0 or P_1 are stationary with respect to a small perturbation $\delta\mathbf{f}$ in that discriminator,

$$\delta P_i = 0, \tag{2.8}$$

either because P_i is constant by design (by making a change δc_T in the threshold if necessary), or because it is extremized with respect to $\mathbf{f}$.

To analyze the changes $\delta P_{0,1}$ further we must know something about the densities $p_i(c)$. If the scanning window is larger than the dimensions of the main visual features to be discriminated (the mine's along and cross range highlight-shadow signature), and if the number of pixels $R \times C$ in the window (2.2) is large, then the random variable c is itself the weighted sum (2.2) of many random variables m_k, and the densities $p_i(c)$ approach the Gaussian distribution by the Central Limit Theorem [12]. The Gaussian assumption is slightly more restrictive than necessary. Here we only require that, as in the Gaussian, the distributions can be parameterized by their respective standard deviation σ_i (the characteristic length scales in Fig.(1)) and mean μ_i (simple translation along the c axis in Fig.(1)); in other words, that we can apply a linear transformation for each class

$$c'_i = \frac{c - \mu_i}{\sigma_i}, \tag{2.9}$$

whereby

$$P_i = \int_{c'_{T_i}}^{\infty} p'_i(c'_i)\, dc'_i, \tag{2.10}$$

with

$$c'_{T_i} = \frac{c_T - \mu_i}{\sigma_i}, \tag{2.11}$$

and

$$p'_i(c'_i) = \sigma_i p_i(c). \tag{2.12}$$

Discrimination in the transform variables c'_i is shown graphically in Fig.(2).

2.2 Prior information

The discriminant mean μ_i and variance σ_i^2 are defined as

$$\mu_i = \langle c \rangle_i = \langle \mathbf{f} \cdot \mathbf{m} \rangle_i = \mathbf{f} \cdot \langle \mathbf{m} \rangle_i, \tag{2.13}$$

and

$$\sigma_i^2 = \left\langle (\mathbf{f} \cdot \mathbf{m} - \mu_i)^2 \right\rangle_i = \mathbf{f}^T \Sigma_i \mathbf{f}, \tag{2.14}$$

in which $\langle \cdots \rangle_i$ denotes the mean taken over the i'th class—by averaging over a large set of actual sonar images representing one particular class for instance; and

$$\Sigma_i = \left\langle (\mathbf{m} - \langle \mathbf{m} \rangle_i)(\mathbf{m} - \langle \mathbf{m} \rangle_i)^T \right\rangle_i \tag{2.15}$$

is likewise the scatter matrix for the i'th class. Hence, given a linear discriminator $\mathbf{f}$, and proceeding on the strength of the Central Limit Theorem, we find that the relevant class information, so far as the performance probabilities are concerned, are 1) the class mean $\langle \mathbf{m} \rangle_i$ in (2.13), and 2) the class scatter matrix Σ_i in (2.14). This is the kind of data that should be compiled and extracted from actual sonar trials, both to determine the performance probabilities for a given linear detector, and to derive the optimal detector as shown below.

In practice, $\langle \mathbf{m} \rangle_i$ and Σ_i may not be available for both the target and clutter classes. For if the seafloor is highly variable throughout a sonar swathe, with one patch of seafloor very different from another, then the clutter class may be poorly represented by a single mean and scatter. Indeed, the mean and scatter for two distinct patches of the seafloor combined into a single class may not be a good representation for either one encountered independently, in which case it may be reasonable to ignore the clutter class altogether, giving the *unapprised* discriminator of the *first kind* listed in Table (1.1). Its complement, the *unapprised* of the *second kind*, assumes the reverse, that nothing is known about the target class while the background clutter is well characterized. This apparently has not been used in practice, but it is an interesting possibility, detecting mines simply because they look different from the background clutter rather than because they look particularly mine like. Reverse discrimination of this kind is feasible when the seafloor clutter is benign and very consistent.

Chapter 3

Unapprised discriminators of the first and second kind

We now derive the mathematical form of the *unapprised* discriminators by optimizing P_1 (first kind) and P_0 (second kind) independently. Assuming a fixed detection threshold c_T, straightforward expansion of equation (2.8) by way of (2.10) gives

$$\begin{aligned} \delta P_i &= -p_i'(c_{T_i}')\,\delta c_{T_i}' \\ &= -\sigma_i p_i(c_T)\left(\tfrac{\partial c_{T_i}'}{\partial \mu_i}\delta\mu_1 + \tfrac{\partial c_{T_i}'}{\partial \sigma_i}\delta\sigma_i\right) \\ &= \sigma_i p_i(c_T)\left(\tfrac{1}{\sigma_i}\delta\mu_i + \tfrac{c_{T_i}'}{\sigma_i}\delta\sigma_i\right) \\ &= p_i(c_T)\left(\delta\mu_i + c_{T_i}'\delta\sigma_i\right) \qquad = 0. \end{aligned} \tag{3.1}$$

The perturbation in the mean and variance of each distribution can be derived from (2.13) to (2.15),

$$\delta\mu_i = \frac{\partial\mu_i}{\partial f_k}\delta f_k = \frac{\partial\left(\mathbf{f}\cdot\langle\mathbf{m}\rangle_i\right)}{\partial f_k}\delta f_k = \langle m_k\rangle_i\,\delta f_k = \langle\mathbf{m}\rangle_i\cdot\delta\mathbf{f}, \tag{3.2}$$

and

$$\delta\sigma_i = \frac{\delta\left(\sigma_i^2\right)}{2\sigma_i} = \frac{1}{2\sigma_i}\frac{\partial\left(\sigma_i^2\right)}{\partial f_k}\delta f_k = \frac{1}{\sigma_i}\delta f_k\left(\Sigma_i\mathbf{f}\right)_k = \frac{1}{\sigma_i}\left(\Sigma_i\mathbf{f}\right)\cdot\delta\mathbf{f}; \tag{3.3}$$

the summation over the repeated subscript k (but not i) being implied. Substituting (3.2) and (3.3) into (3.1) we find that the change in the probability

of detection due to a change $\delta\mathbf{f}$ in the discriminator is

$$\delta P_i = p_i(c_T)\left(\langle\mathbf{m}\rangle_i + \frac{c'_{T_i}}{\sigma_i}\Sigma_i\mathbf{f}\right)\cdot\delta\mathbf{f} = 0, \tag{3.4}$$

which, if true for any perturbation $\delta\mathbf{f}$, implies that

$$\langle\mathbf{m}\rangle_i + \frac{c'_{T_i}}{\sigma_i}\Sigma_i\mathbf{f} = \mathbf{0}, \tag{3.5}$$

giving two optimal solutions

$$\mathbf{f}_i = a_i\Sigma_i^{-1}\langle\mathbf{m}\rangle_i \;\; ; i = 0, 1. \tag{3.6}$$

Here c'_{T_i} and σ_i have been rolled into an arbitrary scale factor a_i—arbitrary because the performance probabilities, being dimensionless, are not affected by a change in the dimension of the discriminator $\mathbf{f}$. The subscript $i = 1$ or 0 for unapprised discriminators of the first and second kind respectively. The discriminators are optimal inasmuch as they maximize the size of their respective non-dimensional thresholds $|c'_{T_i}|$ in Fig.(2) for a given threshold c_T, hence extremizing their respective probabilities P_i: $\mathbf{f}_1$making P_1 as large as possible, or $\mathbf{f}_0$ making P_0 as small as possible.

It is instructive to examine two progressively simplified instances of unapprised discriminators.

3.1 Matched filter

If the scatter is uniform and diagonal

$$\Sigma_i = s_i^2 I, \tag{3.7}$$

where I is the identity matrix, and s_i^2 is a uniform pixel intensity variance, then the optimal unapprised discriminator $\mathbf{f}_i$ in (3.6) is the *matched filter*

$$\mathbf{f}_i = a_i\langle\mathbf{m}\rangle_i. \tag{3.8}$$

The diagonal form (3.7) may in fact be reasonable when the image distortion is due to uniform pixel noise, or when there is no scatter information available. The effect of Σ_i^{-1} in the unapprised discriminator (3.6), then, is to orient the discriminator optimally with respect to the class scatter. The matched filter is also a special case of the *apprised* discriminator considered below (Section (4.2) and (4.3)).

3.2 Mean and energy detector

The simplest discriminator results when we furthermore assume that the relevant class i is very diverse (uniformly random perhaps) such that its mean is simply the uniform vector

$$\langle \mathbf{m} \rangle_i = \overline{m_i}\mathbf{e}, \tag{3.9}$$

where

$$\mathbf{e} = [1, 1, 1, \cdots]^T, \tag{3.10}$$

and

$$\overline{m_i} = \frac{1}{RC} \langle \mathbf{m} \rangle_i \cdot \mathbf{e}, \tag{3.11}$$

is the mean intensity; RC being both the number of pixels scanned in (2.2) and the length of $\mathbf{m}$ and $\mathbf{e}$. The matched filter (3.8) is then

$$\mathbf{f}_i = a_i \overline{m_i}\mathbf{e}, \tag{3.12}$$

and the discriminant (2.5) becomes

$$c = \mathbf{m} \cdot \mathbf{f}_i = a_i \overline{m_i}\ \mathbf{m} \cdot \mathbf{e} = a_i RC\ \overline{m_i}\ \overline{m}, \tag{3.13}$$

in which

$$\overline{m} = \frac{1}{RC}\mathbf{m} \cdot \mathbf{e} \tag{3.14}$$

is the mean pixel intensity of the scanned image. Such a detector would be an energy detector if the dimensions $R \times C$ of the processing window are comparable to the spacial extent of the highlight for a target in the sidescan image.

Chapter 4

Apprised discriminators of the first and second kind

The best performance is expected when given full knowledge ($\langle \mathbf{m} \rangle_i$ and Σ_i) of both classes ($i = 1, 0$). The two performance probabilities P_1 and P_0 are now coupled by prescribing one and simultaneously optimizing the other, which determines both the detection threshold

$$\begin{aligned} c'_{T0} &= \tfrac{c_T(P_0)-\mu_0}{\sigma_0} \quad \text{and} \quad c'_{T1} = \tfrac{c_T(P_0)-\mu_1}{\sigma_1} = \tfrac{c'_{T0}\sigma_0+\mu_0-\mu_1}{\sigma_1} \quad \text{for prescribed } P_0, \\ c'_{T1} &= \tfrac{c_T(P_1)-\mu_1}{\sigma_1} \quad \text{and} \quad c'_{T0} = \tfrac{c_T(P_1)-\mu_0}{\sigma_0} = \tfrac{c'_{T1}\sigma_1+\mu_1-\mu_0}{\sigma_0} \quad \text{for prescribed } P_1; \end{aligned} \tag{4.1}$$

and the discriminator $\mathbf{f}$ in the following way.

Let the prescribed probability be $P_{\bar{\imath}}$ and the optimized probability be P_i to treat both prescriptions in (4.1) at once. The characteristic equation (2.8) for optimization now becomes

$$\begin{aligned} \delta P_i &= -p'_i(c'_{T_i})\,\delta c'_{T_i} \\ &= -\sigma_i p_i(c_T)\left(\tfrac{\partial c'_{T_i}}{\partial \mu_i}\delta\mu_{\bar{\imath}} + \tfrac{\partial c'_{T_i}}{\partial \sigma_i}\delta\sigma_{\bar{\imath}}\right) \\ &= -\sigma_i p_i(c_T)\left[\tfrac{1}{\sigma_i}\left(c'_{T_{\bar{\imath}}}\delta\sigma_{\bar{\imath}} + \delta\mu_{\bar{\imath}} - \delta\mu_i\right) - \tfrac{c'_{T_i}}{\sigma_i}\delta\sigma_i\right] \\ &= -p_i(c_T)\left[c'_{T_{\bar{\imath}}}\delta\sigma_{\bar{\imath}} + \delta\mu_{\bar{\imath}} - \delta\mu_i - c'_{T_i}\delta\sigma_i\right] \qquad = 0. \end{aligned} \tag{4.2}$$

Using (3.2) and (3.3) as before, we find that

$$\begin{aligned} 0 &= c'_{T\underline{i}}\delta\sigma_{\underline{i}} - c'_{Ti}\delta\sigma_i + \delta\mu_{\underline{i}} - \delta\mu_i \\ &= \left\{\left(\frac{c'_{T\underline{i}}}{\sigma_{\underline{i}}}\Sigma_{\underline{i}} - \frac{c'_{Ti}}{\sigma_i}\Sigma_i\right)\mathbf{f} + \left(\langle\mathbf{m}\rangle_{\underline{i}} - \langle\mathbf{m}\rangle_i\right)\right\}\cdot\delta\mathbf{f}, \end{aligned} \tag{4.3}$$

which, being true for all $\delta\mathbf{f}$, implies that

$$\left(\frac{c'_{T\underline{i}}}{\sigma_{\underline{i}}}\Sigma_{\underline{i}} - \frac{c'_{Ti}}{\sigma_i}\Sigma_i\right)\mathbf{f} + \left(\langle\mathbf{m}\rangle_{\underline{i}} - \langle\mathbf{m}\rangle_i\right) = \mathbf{0}, \tag{4.4}$$

whereby

$$\mathbf{f} = a\left(\frac{c'_{T\underline{i}}}{\sigma_{\underline{i}}}\Sigma_{\underline{i}} - \frac{c'_{Ti}}{\sigma_i}\Sigma_i\right)^{-1}\left(\langle\mathbf{m}\rangle_i - \langle\mathbf{m}\rangle_{\underline{i}}\right), \tag{4.5}$$

in which an arbitrary scaling a has again been included. Notice that the optimal discriminator $\mathbf{f}$ is symmetric in i and $\underline{i}$. The apprised discriminators of the first and second kind therefore have the same form, with different coefficients c'_{T0} and c'_{T1} given in (4.1). The solution of (4.5) is complicated by the fact that these coefficients depend on the unknown discriminator $\mathbf{f}$,

$$\frac{c'_{T\underline{i}}}{\sigma_{\underline{i}}} = \frac{c'_{T\underline{i}}\left(P_{\underline{i}}\right)\sqrt{\mathbf{f}\cdot\Sigma_{\underline{i}}\mathbf{f}} + \mathbf{f}\cdot\left(\langle\mathbf{m}\rangle_{\underline{i}} - \langle\mathbf{m}\rangle_i\right)}{\mathbf{f}\cdot\Sigma_i\mathbf{f}} \tag{4.6}$$

$$\frac{c'_{Ti}}{\sigma_i} = \frac{c'_{Ti}\left(P_i\right)}{\sqrt{\mathbf{f}\cdot\Sigma_i\mathbf{f}}}. \tag{4.7}$$

A numerical solution akin to root finding is therefore required.

4.1 Numerical solution

The transformed threshold $c'_{T\underline{i}}\left(P_{\underline{i}}\right)$ is determined by the prescribed probability $P_{\underline{i}}$, and hence by the unknown coefficients $\mathbf{f}$. In practice we must therefore relate $c'_{T\underline{i}}$, $P_{\underline{i}}$, and $\mathbf{f}$ by way of (2.10), by assuming a particular form for the probability density $p'_i\left(c'_i\right)$ in such a way that the dependence on $\mathbf{f}$ comes about solely through the discriminant variance $(\sigma'_i)^2$ and mean μ'_i. This is called a *parametric* approach to the probability density [1]. Figure (3) illustrates the parametric relations for detection ($i = 1$) and false alarm ($i = 0$)

when the Gaussian distribution is assumed in accordance with the Central Limit Theorem as mentioned earlier. We may then compute the optimal discriminator $\mathbf{f}$ numerically by first constructing a family of discriminators

$$\mathbf{f}(\rho) = \left(\Sigma_1 + \rho\Sigma_0\right)^{-1}\left(\langle\mathbf{m}\rangle_0 - \langle\mathbf{m}\rangle_1\right), \tag{4.8}$$

for any scalar ρ, and then plotting the function

$$\begin{aligned} R(\rho) &= -\left(c'_{T_0}\sigma_1\right) / \left(c'_{T_1}\sigma_0\right) \\ &= -\left\{\frac{c'_{T_1}(P_1)\sqrt{\mathbf{f}(\rho)\cdot\Sigma_1\mathbf{f}(\rho)} + \mathbf{f}(\rho)\cdot\left(\langle\mathbf{m}\rangle_1 - \langle\mathbf{m}\rangle_0\right)}{\mathbf{f}(\rho)\cdot\Sigma_0\mathbf{f}(\rho)}\right\}\frac{\sqrt{\mathbf{f}(\rho)\,\Sigma_1\mathbf{f}(\rho)}}{c'_{T_1}(P_1)}. \end{aligned} \tag{4.9}$$

Second, by inspection of (4.5), with allowances for the arbitrary scale factor a, it is clear that optimal discriminator $\mathbf{f} = \mathbf{f}(\rho)$ occurs when ρ satisfies

$$R(\rho) = \rho. \tag{4.10}$$

Note that $R(\rho) > 0$ and $\rho > 0$ because the transformed thresholds c'_{T_0} and c'_{T_1} in (4.1) have opposite signs when c_T lies between μ_0 and μ_1 (as in Fig.(1)).

4.2 Correlation and matched filter

If both scatter matrices Σ_0 have the particularly simple diagonal form (3.7), then the optimal discriminator simplifies to

$$\mathbf{f} = a\left(\langle\mathbf{m}\rangle_1 - \langle\mathbf{m}\rangle_0\right), \tag{4.11}$$

which is sometimes called the correlation discriminator [1]. It is in fact the difference between the matched filters of the first and second kind (3.8), improving on each insofar as $\mathbf{f}$, by analogy with vectors in three-dimensional space, now "points" from the mean of the clutter class 0 to the mean of the target class 1, along the line of optimal discrimination. This makes the discriminator $\mathbf{f}$ insensitive to components of the scanned image lying orthogonal to that "direction", and which therefore do not help to distinguish between the classes. Matched filters, on the other hand, always "point" from the origin to the mean of their respective class, and therefore generally respond in some degree to irrelevant orthogonal components, making the discriminator variance σ_0^2 larger (or equivalently, the relative threshold $|c'_{T_0}|$ smaller) than it need be, degrading performance.

4.3 Improved matched filter performance by de-meaning

The correlation discriminator (4.11) reduces to the matched filter (3.8) if $\langle \mathbf{m} \rangle_0 = \mathbf{0}$. This occurs in practice when the non-mine class is first of all very diverse, such that $\langle \mathbf{m} \rangle_0 = \overline{m_0}\mathbf{e}$ as in (3.9), and when all images are furthermore *de-meaned* before processing—that is, the mean image intensity $\overline{m}$ in (3.14) is subtracted from all images (mine, non-mine, and scanned images alike),

$$\mathbf{m} \leftarrow \mathbf{m} - \overline{m}\mathbf{e}. \tag{4.12}$$

When faced with a very diverse non-mine class, then, the performance of the matched filter can be improved to that of the correlation filter simply by de-meaning before processing.

Chapter 5

Conclusions

It has been shown how the optimal linear discriminator depends on the state of knowledge about the target (mine) and clutter (non-mine) classes to be discriminated. The theory provides practical insight and direction for the object detection problem by

1. identifying the kind of information required for the best possible performance using a linear discriminator—the mean $\langle \mathbf{m} \rangle_i$ and scatter matrix Σ_i for the target ($i = 1$) and clutter ($i = 0$) classes;

2. optimizing the expected performance for a given prior knowledge of each class;

3. ranking the matched filter as one of several optimal linear discriminators, and in this way clarifying the prior knowledge it assumes—a simple scatter matrix for targets, and either total ignorance of the clutter (Section (3.1)), or a very diverse clutter class with zero mean (Section (4.3));

4. and showing when prior local de-meaning of all images can upgrade the matched filter performance to that of a correlation filter (Section (4.3)).

Two points should be made in conclusion. First, that, being probabilistic, the optimization maximizes the performance for many independent trials, but not necessarily for any one trial. Given a sonar image of a mine target range, for instance, one might subjectively tailor a linear discriminator that

is specially suited for the given image, if only by trial-and-error adjustments of discriminator $\mathbf{f}$ and threshold c_T for instance, giving perfect detection with no false alarms for that particular image. But our goal is the best operational performance for the very large population of images expected in practice. The present theory applies operationally in this way because it is based on probability theory. The subjectively tailored discriminator, on the other hand, admits no such generalization.

Secondly, the optimal detector will only be as good as the prior knowledge of the classes used in its design. Thus the optimal discriminator for detecting mines on a rocky seafloor may not be optimal for a sandy, and vice versa. The optimal for both rock and sand considered as a single non-mine or clutter class may furthermore perform worse than that for just rock or sand alone. The role of the clutter class in the optimal discriminator underlines the need for segmentation and classification of the seafloor clutter into distinct subclasses (sand, rock, etc.), and adapting the discriminator to each accordingly.

Research continues on adaptive discriminators of this kind, as well as on compiling representative class means and scatter matrices from actual sonar data, and on assessing various kinds of image preprocessing (compression, de-noising, and non-linear changes of intensity) by their effect on the class means and scatter matrices, and hence on the optimal discriminator and its performance.

Bibliography

[1] K. Fukunaga, *Introduction to statistical pattern recognition*, 2nd Ed., Academic Press, New York, 1990.

[2] L. Devroye, L. Gyorifi, and G. Lugosi, *A probabilistic approach to pattern recognition*, Springer, New York, 1996.

[3] B.D. Ripley, *Pattern recognition and neural networks*, Cambridge university Press, Cambridge, 1996.

[4] H.L. Van Trees, Detection, estimation, and modulation theory, John Wiley and Sons, New York, 1968.

[5] J. Paul, P. Johnson, and J. Harris, "*Final report for the Canadian remote minehunting sonar study*," Canadian National Defence Research and Development Branch Contractor Report DREA CR/96/451A, 1996.

[6] R.T. Kessel, "*Trainable linear filters for the detection of mines in side scan sonar images*," Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 1999-171, Dec. 1999.

[7] G.J. Dobeck, J.C. Hyland, and L. Smedly, "Automated detection/classification of sea mines in sonar imagery," Naval Research Reviews, Office of Naval Research, V.49, pp. 9-20, 1997.

[8] J.C. Hyland and G.J. Dobeck, "Sea mine detection and classification using side-looking sonar," *Proceedings of the SPIE Conference on Detection Technologies for Mines and Minelike Targets*, 17-21 April 1995, Orlando Florida, V. 2496, pp. 466-474.

[9] J. Fawcett, "*Preliminary results of an algorithm for automatic detection of mine-like objects in sidescan sonar images*", Canadian National

Defence Research and Development Branch Technical Memorandum DREA TM 2000-013, Jan. 2000.

[10] R.J. Urick, *Principles of underwater sound for engineers*, McGraw Hill, New York, 1967.

[11] A.D. Waite, *Sonar for practising engineers*, 2nd Ed, Thomson Marconi Sonar Ltd, Great Britain, 1998.

[12] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd edition, McGraw Hill, 1991.

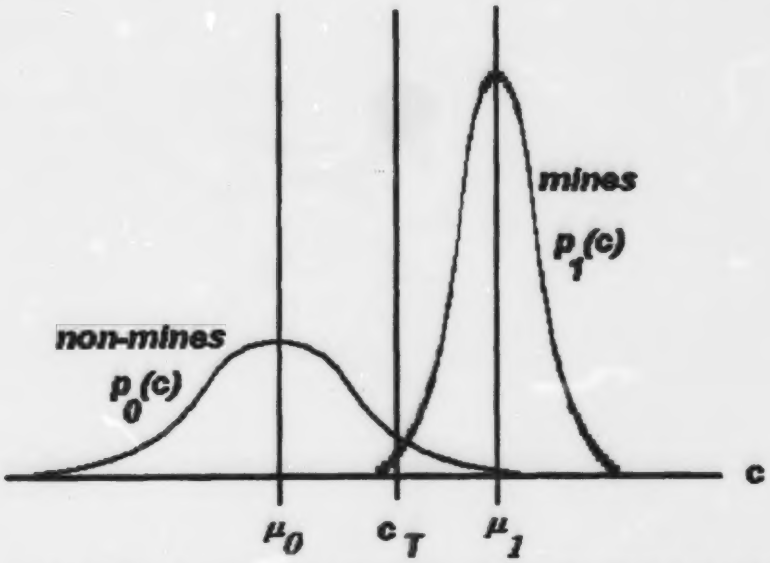

Fig. 1 Conditional probability densities of the discriminant c for non-mines and mines.

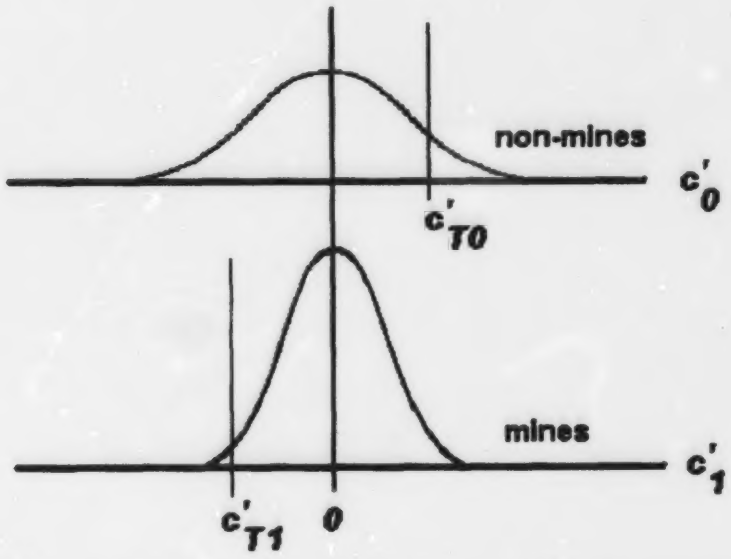

Fig. 2 Conditional probability densities of the transformed discriminant for non-mines and mines.

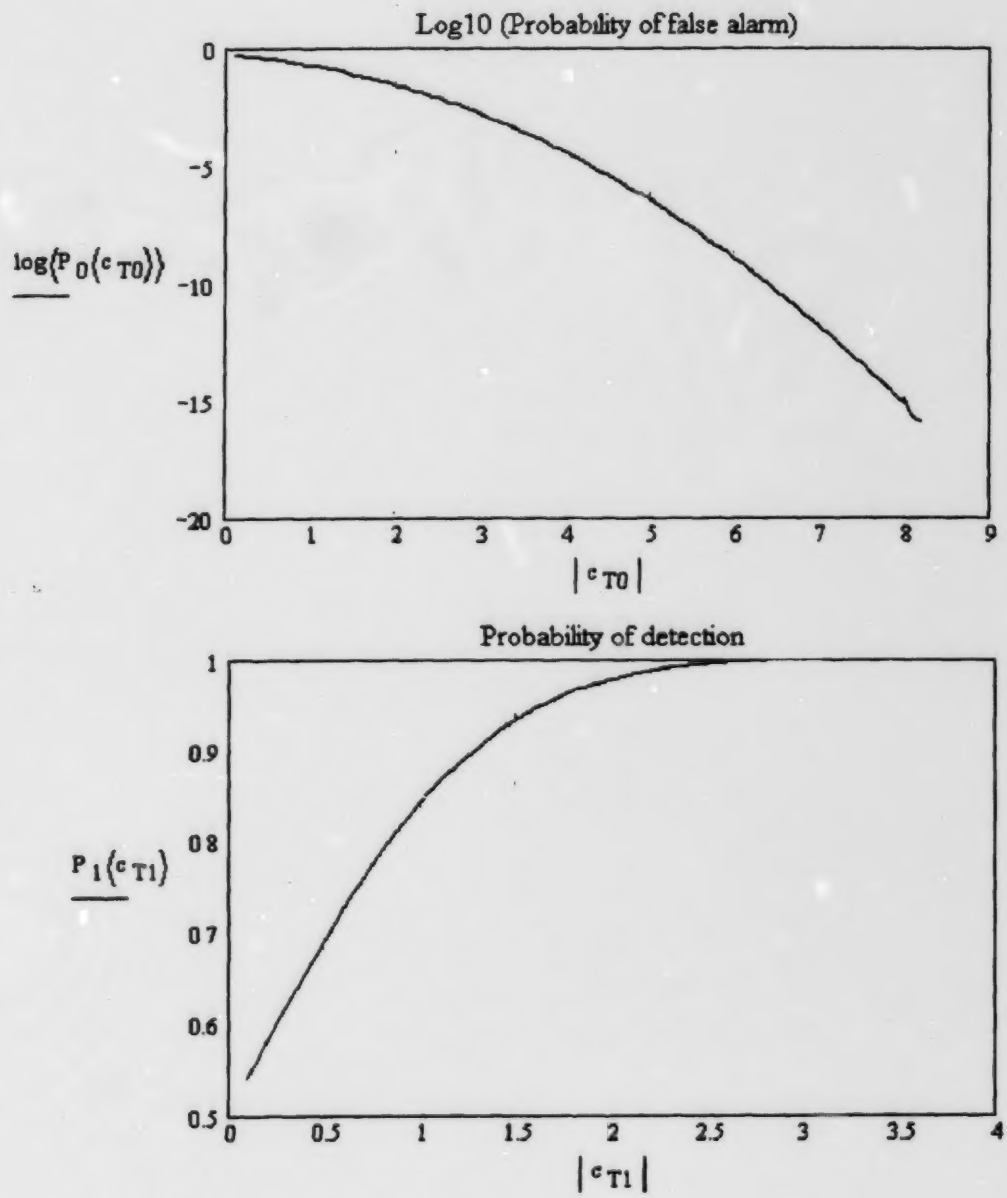

Fig. 3 Conditional probability of false alarm and detection as a function of the magnitude of the transformed detection threshold, assuming Gaussian densities in accordance with the Central Limit Theorem (with zero mean and unit variance as in Fig.(2)).

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1 ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence Research Establishment Atlantic
PO Box 1012
Dartmouth, Nova Scotia, Canada B2Y 3Z7

2 SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

UNCLASSIFIED

3 TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title)

A probabilistic theory for the design of optimal linear discriminators for the automated detection of objects in sidescan sonar images

4 AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Kessel, Ronald T.

5 DATE OF PUBLICATION (month and year of publication of document)	6a NO. OF PAGES (total containing information. Include Annexes, Appendices, etc)	6b NO. OF REFS (total cited in document)
Mar 2000	23 (approx.)	12

7 DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered)

TECHNICAL MEMORANDUM

8 SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address)

Defence Research Establishment Atlantic, Electromagnetics Section, MCM Group

9a PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)	9b CONTRACT NO. (if appropriate, the applicable number under which the document was written).

10a ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b OTHER DOCUMENT NOs (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DREA TM 2000-021	

11 DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12 DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected)

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

13. ABSTRACT (a brief and factual summary of the document It may also appear elsewhere in the body of the document itself It is highly desirable that the abstract of classified documents be unclassified Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U) It is not necessary to include here abstracts in both official languages unless the text is bilingual)

(U) Computerized pattern recognition can be used to help a sonar operator locate underwater objects in sidescan sonar images. The theory behind several linear discriminators is presented here with a view to improving this automated detection. The discriminators are optimal insofar as they maximize the detection performance as defined under the Neyman-Pearson design criteria, with adjustments made to those criteria to suit the prior knowledge of both the objects sought and the local seafloor clutter. The emphasis thoughout is on sea mine detection in naval operations. The theory gives practical insight and direction for the mine detection problem, showing, for instance, 1) what kind of data should be extracted from target and clutter image libraries to get optimal detection performance; 2) that the matched filter, favored for its simplicity, is just one of several optimal linear discriminators resulting in this case when nothing is assumed about the local seafloor clutter; and 3) that prior de-meaning of images will in some cases improve detection performance.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document They should be selected so that no security classification is required Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included If possible keywords should be selected from a published thesaurus e.g Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title)

mine hunting, sidescan sonar, pattern recognition, detection, optimal discriminator

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM